

AD-A071 430

YALE UNIV NEW HAVEN CONN DEPT OF COMPUTER SCIENCE
TEXT PROCESSING EFFECTS AND RECALL MEMORY.(U)

F/6 5/5

MAY 79 W 6 LEHNERT

N00014-75-C-1111

UNCLASSIFIED

RR-157

NL

1 OF 1
AD
A071430



LEVEL ~~II~~

12

A071430



Text Processing Effects and Recall Memory

May 1979

By

Wendy G. Lehnert

Research Report #157

DDC FILE COPY

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DDC
RECEIVED
JUL 19 1979
D

YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

79 07 18 004

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DDC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist.	Avail and/or special
A	

Text Processing Effects and Recall Memory

May 1979

By

Wendy G. Lehnert

Research Report #157

The research described here was done at the Yale Artificial Intelligence Project and is funded in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER #157	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) (6) Text Processing Effects and Recall Memory		5. TYPE OF REPORT & PERIOD COVERED (9) Research Report
7. AUTHOR(s) (10) Wendy G. Lehnert		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University - Department of Computer Science 10 Hillhouse Avenue New Haven, Connecticut 06520		8. CONTRACT OR GRANT NUMBER(s) (15) N00014-75-C-1111
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, Virginia 22217 (12) 56 p1		12. REPORT DATE (11) May 79
		13. NUMBER OF PAGES 49
		15. SECURITY CLASS. (of this report) Unclassified
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this report is unlimited.		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
<div style="border: 1px solid black; padding: 5px; display: inline-block;"> DISTRIBUTION STATEMENT A Approved for public release; Distribution Unlimited </div>		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) (14) RR-157		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Memory Representation Story Understanding Natural Language Processing Dynamic Memory		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) One current research area within the field of natural language processing is concerned with the nature of predictive understanding mechanisms that seek to interpret input in terms of knowledge-based expectations. A predictive understanding system (human or computational) builds an internal representation for narratives by generating inferences and making casual connections between the events described on the basis of these expectations. Progress in this area now makes it possible to study the various ways that processing complexity at		

the time of understanding affects the memory representation generated for a text. For example, memory encodings for the event, "John was shot" will vary depending on whether or not this event was consistent with the understander's expectations (if John was in a duel), inconsistent (if John was watching television), or overshadowed by a strong context (if John was shot along with the President of the United States). In each of these cases, interpretive processing at the time of understanding will vary significantly. Using techniques of artificial intelligence, this project will investigate the ways that processing at the time of understanding can affect recall behavior in the tasks of question answering and paraphrase production. This investigation will focus on the design and implementation of process-oriented memory structures in a computer program that understands narratives. The results of these efforts will be of interest to researchers in cognitive psychology and linguistics as well as artificial intelligence.

-- OFFICIAL DISTRIBUTION LIST --

Defense Documentation Center Cameron Station Alexandria, Virginia 22314	12 copies
Office of Naval Research Information Systems Program Code 437 Arlington, Virginia 22217	2 copies
Advanced Research Projects Agency Cybernetics Technology Office 1400 Wilson Boulevard Arlington, Virginia 22209	3 copies
Office of Naval Research Branch Office - Boston 495 Summer Street Boston, Massachusetts 02210	1 copy
Office of Naval Research Branch Office - Chicago 536 South Clark Street Chicago, Illinois 60615	1 copy
Office of Naval Research Branch Office - Pasadena 1030 East Green Street Pasadena, California 91106	1 copy
Mr. Steven Wong Administrative Contracting Officer New York Area Office 715 Broadway - 5th Floor New York, New York 10003	1 copy
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20375	6 copies
Dr. A.L. Slafkosky Scientific Advisor Commandant of the Marine Corps Code RD-1 Washington, D.C. 20380	1 copy
Office of Naval Research Code 455 Arlington, Virginia 22217	1 copy

Office of Naval Research Code 458 Arlington, Virginia 22217	1 copy
Naval Electronics Laboratory Center Advanced Software Technology Division Code 5200 San Diego, California 92152	1 copy
Mr. E.H. Gleissner Naval Ship Research and Development Computation and Mathematics Department Bethesda, Maryland 20084	1 copy
Captain Grace M. Hopper NAICOM/MIS Planning Board Office of the Chief of Naval Operations Washington, D.C. 20350	1 copy
Mr. Kin B. Thompson Technical Director Information Systems Division OP-91T Office of the Chief of Naval Operations Washington, D.C. 20350	1 copy
Advanced Research Project Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, Virginia 22209	1 copy
Professor Omar Wing Columbia University in the City of New York Department of Electrical Engineering and Computer Science New York, New York 10027	1 copy
Office of Naval Research Assistant Chief for Technology Code 200 Arlington, Virginia 22217	1 copy
Captain Richard L. Martin, USN Commanding Officer USS Francis Marion (LPA-249) FPO New York 09501	1 copy
Major J.P. Pennell Headquarters, Marine Corp. (Attn: Code CCA-40) Washington, D.C. 20380	1 copy

Text Processing Effects and Recall Memory

Wendy G. Lehnert
Computer Science Department
Yale University
New Haven, Connecticut 06520

ABSTRACT

One current research area within the field of natural language processing is concerned with the nature of predictive understanding mechanisms that seek to interpret input in terms of knowledge-based expectations. A predictive understanding system (human or computational) builds an internal representation for narratives by generating inferences and making casual connections between the events described on the basis of these expectations. Progress in this area now makes it possible to study the various ways that processing complexity at the time of understanding affects the memory representation generated for a text. For example, memory encodings for the event, "John was shot" will vary depending on whether or not this event was consistent with the understander's expectations (if John was in a duel), inconsistent (if John was watching television), or overshadowed by a strong context (if John was shot along with the President of the United States). In each of these cases, interpretive processing at the time of understanding will vary significantly. Using techniques of artificial intelligence, this project will investigate the ways that processing at the time of understanding can affect recall behavior in the tasks of question answering and paraphrase production. This investigation will focus on the design and implementation of process-oriented memory structures in a computer program that understands narratives. The results of these efforts will be of interest to researchers in cognitive psychology and linguistics as well as artificial intelligence.

1. Background to the Problem

The problem of text comprehension is traditionally partitioned into three phases: (1) initial understanding or encoding, (2) retention or memory maintenance, and (3) retrieval for the purposes of answering a question, producing a paraphrase, or performing some other demonstrative task. A successful demonstration of comprehension relies on competence at each stage; an error at any stage can manifest itself in the retrieval task, although it is not always clear, for any given error, exactly where the processing has broken down. Natural language processing research in artificial intelligence has produced process models for text comprehension which concentrate primarily on the initial encoding [11,15,19,20,22,28,29,30,31] and retrieval phases [12,13,29,31]; very little attention has been given to the retention phase.

Standard strategies for the initial encoding phase include syntactic parsers [11,15,29,30], conceptual analyzers [16,18,19,25], predictive knowledge structures [2,6,7,9,24,27], and story grammars [21,26]. Much of the work in this area is very recent, but already there is psychological evidence substantiating the strategy of conceptual analysis guided by predictive knowledge structures [4]. Further results with free recall experiments have argued in favor of semantically-oriented memory representations of the sort generated by knowledge-based systems while arguing against the story grammar approach [5]. Within this knowledge-based viewpoint the recognition of valid causalities between events is thought to be a central factor affecting memory retention [5,23].

Various knowledge structures have been proposed for the purpose of text comprehension. Minsky's system of frames [16] describes a strategy for knowledge application on a very abstract level, while more concrete formulations of knowledge structures have been investigated via computer implementations. The SAM system was designed to investigate the process of script application [9], and PAM was implemented to illustrate how goal-oriented analysis aids understanding [27]. Scripts and goals are two types of knowledge structures proposed by Schank and Abelson [24].

The aspect of knowledge-based understanding we will address here concerns the ways in which encoding processes directly affect retention and retrieval. For example, the von Restorff effect has long been acknowledged in psychology: strange or inherently interesting events within a narrative will be retained for long periods of time and recovered in free recall [8]. But a precise mechanism responsible for this general phenomenon has yet to be proposed. Exactly how is a system going to recognize when an event is strange or inherently interesting? It is not enough to know if an event is causally coherent in a given text because some events are both irrelevant (not causally connected) and boring (neither strange nor interesting). It appears that some further explication of this phenomenon can now be proposed in view of recent research on predictive understanding systems.

We also know that cued recall will elicit more information from memory than free recall [3]. In terms of text processing, this means that a person answering questions about a story will be able to

demonstrate retention of information that will not be included in a free recall paraphrase of the same story. Depending on the nature of the retrieval task, some information will be more accessible than others. While this initially appears to be no more than a problem in retrieval, we believe that encoding processes are critical in determining the variable accessibility of information in memory.

In this paper we will outline a strategy for dynamic memory representation. Dynamic memory is intended to augment standard memory representations with information about processing that occurred at the time of understanding. For example, if we have a memory representation encoding the fact that John hit Mary, the dynamic aspect of this information will tell us exactly how we understood the event when we first heard about it. Was this event unexpected or was it consistent with knowledge-based predictions? Did it contradict specific expectations? Was it overshadowed by surrounding events of greater significance? The answers to these questions will determine how we store the fact that John hit Mary in memory, and thereby affect the ways in which we access this event. We will propose four structures for dynamic memory encoding and show how they can be used to predict various recall behavior in both question answering and paraphrase tasks.

2. Dynamic Memory Structures

The four dynamic memory structures we will propose are defined in terms of knowledge-application processes within a predictive understanding system. A number of computer programs have been designed to investigate problems specific to knowledge structures and

their application in text understanding [2,6,7,9,18,27]. We will limit ourselves to the knowledge structures developed by Roger Schank and Robert Abelson [24]:

A GOAL of X is a state that X desires. Once we know that X has a specific goal we can make predictions about what X might do. These predictions help us to "understand" and interpret X's behavior appropriately. Some examples of goals are "satisfy-hunger," "preserve-health," and "achieve-status."

A PLAN is a general strategy for attaining a goal state. Plans are invoked by specific goals. Goal-oriented behavior is frequently understood in terms of plans that are instrumental to the achievement of that goal. Some examples of plans are "delta-prox" (change of location), "delta-cont" (change of possession), and "delta-soccont" (change of social control).

A SCRIPT is a stereotypic event sequence in a specific situational context. Scripts describe cultural conventions and serve to connect events when a low-level causal analysis of the situation might fail. Examples of scripts include "\$restaurant" (eating in a restaurant), "\$telephone" (placing/receiving phone calls), and "\$gun" (shooting a gun).

These three knowledge structures are used predictively, to anticipate events that we are liable to hear about in the context of a story. The predictive application of these structures also serves to provide

the system with inferences (assumptions that could be wrong) that are critical to overall comprehension. For example, if we are told, "Mary was hungry. She got into her car..." we should predict that Mary is going out to either buy food or to eat. This prediction relies on goals (s-hunger), plans (d-prox), and scripts (\$car, \$store, \$restaurant). If we are then told, "She drove to a drive-in," we will assume that the "drive-in" refers to a fast food restaurant, not a drive-in movie. This interpretation is prediction-driven, relying on a knowledge-based understanding of Mary's goals and behavior.

If an understanding system has access to scripts, plans, and goals, any given input might be understood by script application (with a script-based prediction), plan application (with a plan-based prediction), or goal analysis (with a goal-based prediction). At the same time, an event may fail to be understood by a predictive structure, or more importantly, events can contradict knowledge-based predictions. The encoding process that preserves input information in memory must also tell us precisely which processes received the input information and how that input came to be integrated into the story representation. This process-oriented part of memory will be called dynamic memory. Within the context of an understanding system that exploits predictive knowledge structures, we can define four dynamic memory structures:

[1] PREDICTIVE RESIDUE

Indicates that an event was

- 1) Predicted at the time of understanding, and
- 2) Subsequently substantiated by inference or explicit input

[2] ATTENTION TRANSFER MARKER

Indicates where attention shifted to

- 1) An unpredicted event, or
- 2) A predicted event of inherent interest

[3] EXPECTATIONAL TRACE

Preserves a set of expectations which were

- 1) Alive at one point during understanding, and
- 2) Subsequently contradicted by explicit input

[4] MEMORY EXPANSION MARKER

Points to an instantiation kernel for a script expansion

Each of these structures indicates what distinct type of processing occurred at the time of understanding in order to integrate new input into the story representation. Processing complexity at the time of understanding affects the memory representation generated for a text. A process-oriented analysis of text will therefore help us to predict

various recall phenomena associated with question answering and paraphrase tasks.

Each of the four dynamic memory structures is expected to affect text recall. The precise relationships between dynamic memory structures and recall phenomena require further investigation; here we will outline a few of the areas which appear to be the most promising.

3. Predictive Residues and Question Answering

When people answer questions about stories, they are able to interpret those questions according to the context in which they are asked. In other words, questions are understood by processes that are sensitive to the story at hand. For example, consider the following story:

John closed up his office and went to an office party after work on Friday. While he was there, he overheard a conversation between two executives concerning a special account that John had been working on. From what they were saying, it became clear that John was not receiving credit for the time and energy he had devoted to the project. After hearing this, John felt extremely frustrated and in no mood to socialize. In an effort to get his mind off his troubles, John excused himself and went out to a movie. After the movie he went out and got drunk.

After reading this story, one could be asked,

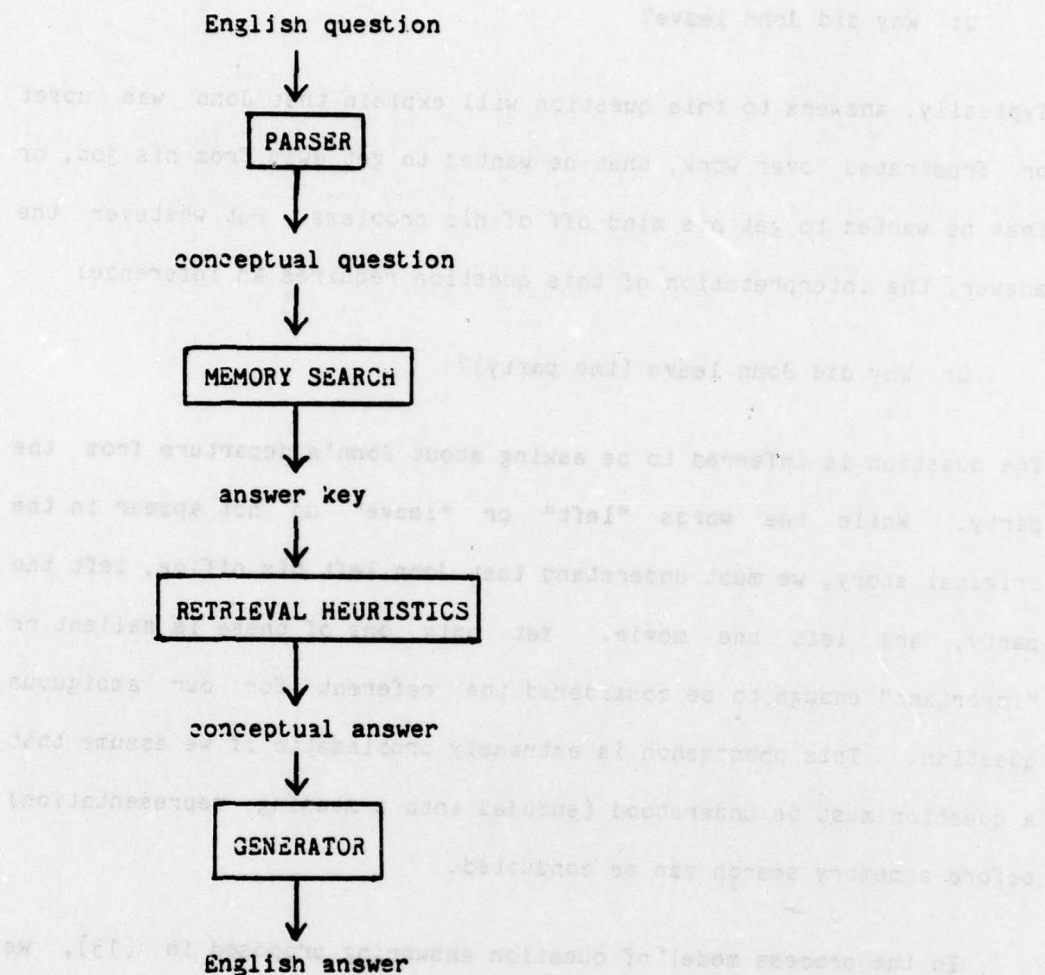
Q: Why did John leave?

Typically, answers to this question will explain that John was upset or frustrated over work, that he wanted to get away from his job, or that he wanted to get his mind off of his problems. But whatever the answer, the interpretation of this question requires an inference:

Q: Why did John leave (the party)?

The question is inferred to be asking about John's departure from the party. While the words "left" or "leave" do not appear in the original story, we must understand that John left his office, left the party, and left the movie. Yet only one of these is salient or "important" enough to be considered the referent for our ambiguous question. This phenomenon is extremely problematic if we assume that a question must be understood (encoded into a meaning representation) before a memory search can be conducted.

In the process model of question answering proposed in [13], we have the following flow of control:



The computer implementation for this process model was insensitive to context in the sense that the parsing procedures were not influenced by information from the story representation. Questions were parsed without any recourse to contextual factors, in much the same way that MARGIE [22] parsed sentences in a contextual vacuum.

We first acknowledged a problem with this state of affairs in [13] while discussing the problem of focus assignments for questions. It was observed that lexically identical questions can assume distinct focus assignments which are dependent on story contexts. For example,

consider the following story:

John had just bought a new car. He was so happy with it that he drove it at every possible opportunity. So last night when he decided to go out for dinner, he drove over to Leone's. When he got there he had to wait for a table....

After completing this story, we can answer the question:

Q: Why did John drive to Leone's?

Answers for this question will explain about John's new car and his desire to drive it. But now suppose we are told a slightly different story:

John had a crush on Mary. But he was so shy that he was happy to merely be in her proximity. So he was in the habit of following her around a lot. He knew that she ate at Leone's very often. So last night when he decided to go out for dinner, he drove over to Leone's. When he got there he had to wait for a table....

Now suppose we are asked the identical question:

Q: Why did John drive to Leone's?

Now our answers will explain that Mary was going to Leone's and John wanted to be near her. As these different answers indicate, our conceptual interpretation of this question varies with each story. Within the context of the first story, the question is assumed to be asking about John's method of transportation. Why did he drive

instead of walk or take a bus or use some other method of vehicular transportation? Within the context of the second story, the question is assumed to be asking about John's destination. Why did he go to Leone's as opposed to a hamburger stand or some other place to eat? The question focusses on driving in the first case, and Leone's in the second.

Intuitively, we would like to say that the assignment of focus for these questions is part of the understanding process; that is, we understand these two questions to be asking about two different things. But where does the focus assignment come from? The conceptual focus for these questions relies on information we have in memory about John's motives for driving to Leone's. If the parser is going to place a focus on the conceptual question that it produces, the parser must in some way be able to access or respond to information in the story representation.

Alternatively, in a less intuitive approach, we can argue that the conceptual question remains ambiguous with respect to focus until the story representation can be examined during the memory search. In other words, we cannot assign a focus to these questions until we have searched memory, and found an answer for the question. The answer then allows us to construct a focus which is consistent with that answer. If we have information in memory explaining "drive," we place the focus on drive; if we have information explaining "Leone's," we focus on Leone's. Any subjective sense of focus is then established as sort of a side effect from the answer found in memory. While this possibility was not ruled out in our earlier discussions, we would now

like to argue against it on the grounds that it forces focus into a completely extraneous position. Why generate a focus for the question unless that focus is in some way useful in finding an answer? Why would an efficient information processing scheme produce information which served no apparent purpose? If we can design an alternative process model which exploits focus, such an alternative must be the preferred model. While the process model in [13] afforded us no such alternative explanation of focus assignment, we can now argue for the exploitation of focus by using predictive residues in dynamic memory.

Predictive residues will enable us to guide the process of understanding so that questions will be interpreted in a manner which is sensitive to the context in which they are asked. Whenever a predicted event is substantiated, that prediction is saved within a predictive residue for its corresponding event. In the course of reading a story, a multitude of predictive residues are generated, each preserving a prediction along with its associated event in memory. These predictions are arranged into a hierarchical structure at the time of understanding, to later compete for the interpretation of questions about that story. In our sample story, there is a prediction for each of the three departures, but only one of these predictions involves high level goal analysis. John's goal of recognition (achieve-status) is frustrated and we expect a new plan (e.g., John decides to work harder), a goal substitution (e.g., John might look for a new job), or an emotional withdrawal (e.g., John sulks). Leaving the party qualifies as a form of withdrawal, and so we understand this departure in terms of a high level goal analysis. Leaving the movie, on the other hand, is merely a script-based

prediction. The hierarchy of predictive residues is constructed to reflect these differences, so John's departure from the party is ranked above leaving his office and leaving the movie. This higher priority is responsible for the interpretation of our question as an "obvious" reference to leaving the party. In this way, predictive residues allow us to interpret technically ambiguous questions in a natural manner, within the context of the story in question.

The ambiguous focus problem is also handled by the appropriate predictions for a given context. In the story about John's new car, we have a strong prediction about John driving his car. In the story about John's infatuation, we have a strong prediction about John going wherever Mary goes. The question "Why did John drive to Leone's?" will satisfy either of these predictions because it tells us that John was driving as well as that John went to the place where Mary was. Our sense of focus in these questions is a function of those conceptual components in the question which match the predictive residue at hand.

In addition to the interpretive advantages of predictive residues, there is another extremely important advantage they provide in terms of the memory search needed to find an answer. In QUALM [12,13], all answers relied on the identification of an answer key in memory. An answer key was that conceptual description in the story representation that corresponded to the question's conceptual content. For example, the question "Why did John leave New York?" would be answered by first locating an answer key in memory encoding the fact that John left New York.


In order to find answer keys, QUALM conducted a brute force search of the story representation using the question concept and a pattern matcher. When a concept in memory was found matching the question concept, we had located the answer key. This brute force search predicts two unfortunate reaction time phenomena: (1) That questions about events occurring at the end of a story will take longer to answer than questions about events at the beginning. (This is assuming that the search initiates at the beginning of the story representation). (2) The longer the story, the longer it will take to answer questions about events near the end. Neither of these predictions seem right, but they are nevertheless unavoidable if a brute-force pattern-matching memory search is conducted. An indexing scheme could lessen the effect (from a linear response to a logarithmic response) but the effect would still exist.

Here we see another important ramification of predictive residues. If a question is interpreted by means of a predictive residue, we can transcend the entire problem of a memory search by merely exploiting a link from the residue to its corresponding data entry in the standard story representation. This link is trivial to generate at the time of understanding, since a substantiated expectation builds its data entry at the time it is satisfied. We need only establish a pointer from that expectation to its data entry in the standard story representation. Given such a link, we can automatically locate the answer key without any search whatsoever. Answer keys are then found in constant time after question interpretation. In other words, it takes no longer to find the answer key in memory than it does to produce a conceptual interpretation of


the question.

The identification of appropriate answer keys becomes even more problematic when we rely on an answer key for more complex question answering behavior. For example, in the SAM system, we implemented a heuristic that would provide elaborations in the event that a verification question (yes or no question) were answered negatively. For example, if SAM were asked "Did the waitress give John a menu?" it could return "No, the hostess gave John a menu." But this elaboration must be derived from a partial match to the question concept. That is, an answer key must be found in the story representation which corresponds closely, but not exactly to the question concept. In this case, the correspondance breaks down with the slot fillers in the "actor" and "from" slots:

QUESTION CONCEPT:




DESIRED ANSWER KEY:



But exactly how do we know to let the actor and from slot fillers vary? If we allow partial matches with other slots, we will get different elaborations:

ALTERNATIVE ANSWER KEY:



If we allow for a variation in the object slot filler, we would get "No, the waitress gave John a check." While this is a technically correct response, it is somehow less appropriate than a response which tells us who gave John a menu. In the SAM system a script-based focus heuristic was implemented to handle the problem of partial pattern matches for answer elaborations. This heuristic relied on the notion of "constant" events in a script and "variable" components within those events. So, for example, a script constant in the restaurant script describes the patron receiving a menu. Where he gets it is a variable: it might come from a waiter, waitress, hostess, maitre d', or he may pick it up off the table himself. The search for an answer key in these cases must then be guided by the script constants and variables.

While this strategy sufficed for the script-based texts processed by SAM, it did not have an immediate analogue in domains dependent on plans and goals. For example, suppose we had read:

A small twin engine airplane carrying federal marshals and a convicted murderer who was being transported to Leavenworth crashed during an emergency landing at O'Hare Airport yesterday. During rescue operations, the murderer was able to grab a gun from a wounded marshal and effect an escape in a fire department vehicle.

In the context of this story it would be reasonable to ask:


Q: Did the murderer get a knife?

To which we would like to respond:


A: No, the murderer got a gun.

Again we have the problem of a partial pattern match, where the wrong match will produce an inappropriate elaboration.


QUESTION CONCEPT:



DESIRED ANSWER KEY:



ALTERNATIVE ANSWER KEY:



Here we see that it is no longer enough to constrain the partial match in terms of slots alone; we now appear to need constraints on the slot fillers as well. In this case we are looking for an ATRANS of a weapon. But where do these constraints come from? What we really need here is knowledge about the murderer's general plan to procure a

weapon. Such an expectation is generated at the time of understanding, distinct from another expectation about procuring a vehicle. If predictive residues can be accessed to aid in the understanding of the question, then this question will be picked up by the residue involving a weapon.

In exploiting predictive residues for memory searches, we see how the intuitive distinction between interpretation and memory retrieval becomes somewhat blurred. Predictive residues are primarily an interpretive mechanism. But once a residue-driven interpretation has been produced, the responsible residue can be further exploited to show us precisely where in the story representation we must look to find the answer. This technique allows us to virtually parse "into the answer key" as well as into a conceptual representation for the question. Predictive residues therefore augment standard story representations with a "top-down" component that guides the interpretation of questions into that story representation. The parsing of questions can then be followed directly with procedures for finding an answer; no intermediate search is needed in order to tie the question in with the story.

Predictive residues are therefore utilized for the task of question answering in two ways: as an interpretive device and as an alternative to massive memory searches. Of the four dynamic memory structures, predictive residues are most useful for the task of question answering. The remaining three dynamic memory structures are used in the task of paraphrase rather than question answering. But before describing these, we will first motivate our discussion by

presenting a sample narrative along with a simplified version of its memory representation.

4. "The War of the Ghosts"

In order to illustrate our three remaining dynamic memory structures we will consider a frequently referenced narrative that Frederic Bartlett first discussed in his early work on human memory [1]. "The War of the Ghosts" is an Eskimo folk tale which Bartlett used to test free recall. Subjects were asked to read the story and then produce a paraphrase of it. His discussion of his experiments is particularly useful for our purposes because he included a number of sample paraphrases taken at various times after the initial reading. We will periodically reference these results in the discussion to follow. But first, the story itself:

The War of the Ghosts

One night two young men from Egulac went down to the river to hunt seals, and while they were there it became foggy and calm. Then they heard war-cries, and they thought: "Maybe this is a war-party". They escaped to the shore, and hid behind a log. Now canoes came up, and they heard the noise of paddles, and saw one canoe coming up to them. There were five men in the canoe, and they said:

"What do you think? We wish to take you along. We are going up the river to make war on the people".

One of the young men said: "I have no arrows".

"Arrows are in the canoe", they said.

"I will not go along. I might be killed. My relatives do not know where I have gone. But you", he said, turning to the other, "may go with them."

So one of the young men went, but the other returned none.

And the warriors went on up the river to a town on the other side of Kalana. The people came down to the water, and they began to fight, and many were killed. But presently the young man heard one of the warriors say: "Quick, let us go home: that Indian has been hit". Now he thought: "Oh, they are ghosts". He did not feel sick, but they said he had been shot.

So the canoes went back to Egulac, and the young man went ashore to his house, and made a fire. And he told everybody and said: "Behold I accompanied the ghosts, and we went to fight. Many of our fellows were killed, and many of those who attacked us were killed. They said I was hit, and I did not feel sick".

He told it all, and then he became quiet. When the sun rose he fell down. Something black came out of his mouth. His face became contorted. The people jumped up and cried.

He was dead.

Using techniques of predictive understanding, an understanding system would internally represent this story on four representational levels: the story includes (1) scripts, (2) plans, (3) goals, and (4) specific events linking these larger structures. We will not attempt to present a complete story representation here, but we will look at a somewhat simplified version which will suffice for our purposes. In particular, we will not try to represent anything connected with the problematic presence of ghosts.

MEM1: script = \$hunt

nunters = 2 men (hereafter referred to as M1,M2)

animal = seal

path = default with ATM1 interrupt

ATM1: nunters <=> ATTEND <- ears (to) MOBJ: war cries



Initiate/Reason

MEM2: script = \$war

attackees = M1,M2



Initiate

P-Healthn goal (Goal1) on part of M1, M2



Reason

M1,M2 <=> PTRANS <- M1,M2 (to) behind-log

(Plan1 in response to Goal1)

MEM2: canoe <=> PTRANS <- canoe (to) M1,M2



enable

conversation between warriors _M1,M2:

This story representation is lacking many details included in the original narrative. These omissions are not intended to reflect assumptions we would like to make about which information is remembered or forgotten. We merely wanted to present a somewhat simplified version of the story representation in order to illustrate various aspects of dynamic memory structures.

MEM1 represents an instantiation of the hunting script which specifies a typical hunt with M1 and M2 acting as hunters and seals acting as their prey. This otherwise uneventful episode is interrupted by ATM1.

ATM1 describes the hunters hearing war cries. This event initiates a reasoning process by M1 and M2 which leads them to conclude that they are part of a new scriptal situation: MEM2.

MEM2 represents an instantiation of the war script which specifies only that M1 and M2 are being attacked by unknown assailants.

The realization of MEM2 initiates a new goal for M1 and M2 which specifies the preservation of their (good) health states. This goal implicitly suggests that their health states are being threatened: M1 and M2 are in danger.

The health goal explains why M1 and M2 hide behind a log. This transfer of location is part of a plan instrumental to health preservation.

MEM2 resumes with the arrival of a canoe filled with warriors.

The arrival of the canoe enables a conversation.

conversation between warriors & M1,M2:

W's: express desire for M1,M2 to join attackers in \$war	Goal2: M1,M2 join #attackers Plan2: Inform/Reason planbox	P1
M1: explain missing enablement (no arrows)	Goal3: Foil Goal2 Plan3: Inform/Reason	P2
W's: explain enablement OK (arrows in canoe)	Goal4: Thwart Plan3 Plan4: Inform/Reason	P3
M1: explicit refusal "relatives don't know..." offer M2 instead	Goal3: Foil Goal2 Plan5: Inform/Reason	P4

```

////////////////////////////////////
/                               /
/           Goal Fate Graph:   /
/                               /
/           P1         P2         P3         P4       /
/                               /
/   Goal1   -----> S         /
/   Goal2   -----> S         /
/   Goal3   -----> TF -----> S       /
/   Goal4   -----> S         /
/                               /
/   S = success   TF = temporary failure /
/                               /
/   The success of Goal1 is only success for M1. /
/                               /
/   If we examine the fate of the competing goals /
/   Goal1 and Goal2, we see that a compromise /
/   was reached satisfying both parties. /
/                               /
////////////////////////////////////

```

The conversation between M1 and the warriors can be analyzed in terms of conflicting goals and their corresponding plans:

GOAL1 (preserve health) is the top-level goal of M1

GOAL2 is the warriors' top-level goal of getting M1 and M2 to join them in the aggressor role of the war script. To achieve this goal they employ a simple inform-reason plan of expressing their desire.

GOAL3 is a subgoal of GOAL1 which opposes the achievement of GOAL2.

GOAL4 is a subgoal of GOAL2 which opposes the achievement of GOAL3.

When GOAL4 succeeds, GOAL3 temporarily fails. PLAN3 is then replaced by PLAN5 which offers a compromise designed to satisfy GOAL2 in part.

Acceptance of the compromise signals the achievement of GOAL3 and its parent GOAL1, along with the warrior's top level GOAL2.

MEM1: \$hunt ends with

M1 <=> PTRANS M1 (to) home, and

M2 joins attackers in \$war

MEM2: (cont. of \$war)

attackers <=> PTRANS <- attackers (to) town

\$fighting

\$death (multiple instances)

ATM2: \$hit with M2 = victim



Initiate/Reason

warrior <=> SPEAK



Initiate

ATM3: M2 <=> MBUILD <- [ghosts?]

ET1: ATM2 does NOT result in pain for M1

MEM2: \$war ends with

attackers <=> PTRANS attackers (to) home



enable

MEM3: \$build-fire with actor = M2

gravity <=> PROPEL <- M2 (to) ground

ATM4: force <=> PROPEL <- black (from) mouth of M2

physstate M2 becomes -10 ..



Initiate/Reason

people <=> SPEAK

ET1 becomes
a PR

MEM1: At this point we are told that M1 goes home and M2 joins the war party, from which we must conclude that \$hunt is no longer an active script.

MEM2: With M2 added to the war party, the \$war script is now continued when the attacking party arrives at a town across from Kalana.

Subscripts within \$war refer to fighting and deaths.

ATM2: In particular, M2 is hit by an arrow. This is worthy of attention because it threatens the life of a main character. Expectations are generated at this point about M2's health state.

M2 getting hit motivates the warrior's observation.

ATM3: The warrior's observation leads M2 to conclude something about ghosts. This conclusion is worthy of attention because it makes no sense and it involves ghosts which are inherently interesting.

ET1: The expectations generated at ATM2 are explicitly contradicted.

MEM2: The war party returns home from which we must conclude that the \$war script is no longer active.

MEM3: a new script \$build-fire is instantiated with a default path specification and a role binding for M2.

M2 tells his story to the villagers.

M2 falls down.

ATM4: Something black (???) comes out of M2's mouth.

M2 dies. At this point ET1 becomes a predictive residue.

The villagers react.

5. Attention Transfer Markers and Free Recall

The notion of an attention shift is critical for the integration and subsequent retention of information in memory. Consider the following excerpt from the "The War of the Ghosts":

"...When the sun rose he fell down. Something black came out of his mouth. His face became contorted. The people jumped up and cried. He was dead."

In all of the paraphrases for this story that Bartlett published [1], subjects mentioned black coming out of his mouth. This event is memorable because it is so unexpected and inexplicable. If we had been told that blood came out of his mouth, we would expect less recall for this event since it can be understood in terms of common knowledge about violent deaths. The fact that it was something black that came out of his mouth is very close to making sense but it does not quite fit (can blood be black?); this serves to enhance its chances of being retained and subsequently reproduced in free recall. Attention transfer markers are designed to reflect the processing history for input information in a way which effectively tells us whether or not we have striking information.

While it is not clear exactly what should be included in an ATM, we can expect that attention transfer markers will be given weights relative to the surrounding context. These relative weights will reflect the ways that what is significant or interesting in one story may be routine or less deserving of attention in another. For example, if a variation on "The War of the Ghosts" were to end with:

..When the sun rose he fell down. Something black came out of his mouth. His face became contorted. A great snake materialized over his body. The snake sang an eerie song and vanished in the air. The people jumped up and cried. He was dead."

it is likely that the singing snake will overpower everything else at this point in the story representation. We could expect subjects to remember the snake, and perhaps do so at the expense of forgetting the black expulsion. If too many interesting or bizarre things happen, some will be designated as lesser events in order to focus attention on a smaller number of more demanding events. We expect there is an upper bound on the ratio of total attention strengths to input processed (TAS/IP ratio) that cannot be surpassed in the course of constructing a story representation. Processing resources are finite, and the amount of processing that can be devoted to problematic input is limited by the resources available [3]. So if the TAS/IP ratio is exceeded, attention transfer marker strengths must be devalued appropriately. An "ATM overload" phenomenon can be expected to occur when there are too many ATM's. Under these conditions we can expect free recall to suffer from a number of omissions. For example, anyone reading the following story would probably suffer from an ATM overload:

John was walking down the street eating an ice cream cone. He saw a man walk into the bushes and begin to undress. Soon a crowd had gathered and the police came to investigate. While they were there a giant explosion

occurred two blocks away. People came running in their direction screaming that there had been a terrible accident. Many were bleeding and one man had lost an arm. Meanwhile a fire broke out in the park. People said there was a conspiracy afoot because a bomb had been sighted nearby only yesterday. When an epidemic broke out the following week, everyone knew the aliens had landed.

A boring story is one for which the TAS/IP ratio fails to exceed some minimal threshold. Stories of the sort understood by SAM [9] which are entirely script-based fall into this category. Optimal retention will presumably occur when the TAS/IP ratio rests within some preferred interval.

There may be a number of secondary memory effects associated with ATM's and input information surrounding them. For example, there may be a retroactive inhibition effect on those concepts immediately preceding an ATM. This would parallel an effect that occurs in list learning experiments where a "surprising" or "unexpected" list element will interfere with the retention of list elements immediately preceding that element. Of course, in the task of text comprehension, we should expect a number of factors to contribute to the reliability of any such phenomenon. In particular, we must assume that the conceptual content of information preceding an ATM will play a major role in whether or not that information will be lost in a free recall account of the story. Information that is necessary for a critical causal connection should be less expendable than information with no causal implications.

In terms of paraphrase behavior, we expect all ATM's to be included in free recall accounts for a story. To see how this prediction fares in Bartlett's study, we check for references to the four ATM's in "The War of the Ghosts" in ten paraphrases Bartlett provided. The results indicate that our ATM's do tend to be preserved in paraphrases, although some are more prevalent than others:

ATM1 (MEM1 interrupt)	.7
ATM2 (M2 is ait)	.9
ATM3 (MBUILD "ghosts")	.6
ATM4 (black from mouth)	1.0

In our dynamic memory model, we would like to generate weights for ATM's that predict the probability of their free recall. According to these results, ATM2 and ATM4 are the strongest ATMs while ATM1 and ATM3 are the weakest. The two weakest ATM's here (ATM1 and ATM3) are interesting because there are additional factors affecting our memory representation at these points. In every case where ATM1 is omitted, it was also the case that the entire script (MEM1) was also omitted from recall. Whenever MEM1 was present in a paraphrase, ATM1 was also present. Since a script interruption cannot be recalled without referencing that script, this result is not surprising. If we had identified ATM1 more generally as merely a script interruption (for any introductory script at the beginning of the story) then ATM1 would have occurred in 3 of the 10 paraphrases.

The contrast between ATM3 and ATM4 is of special interest because they both fall into the category of inexplicable events. Yet ATM3 was included in only 6 of the 10 paraphrases while ATM4 was included in

all 10. One possible explanation derives from the fact that ATM4 describes a concrete event that can be visualized, while ATM3 describes a cognitive event which does not yield a graphic image. Further evidence for this explanation can be found in the fact that ghosts (which can be visualized) are mentioned in some way in 9 of the 10 paraphrases. We might therefore conclude that graphic concepts should receive heavier ATM weightings than abstract concepts.

We expect two major factors to influence the weight of an ATM: (1) shifts in the applicability of various knowledge structures, along with (2) sensitivity to the notion of "inherent interest." Significant factors of the first kind include script triggerings, script interference points, script interruptions, goal activation, plan activation, role theme activation, goal frustration, and goal conflicts. Factors of the second sort will involve a taxonomy of interesting events which must specify instances of violence, sexual activity, and various attributes that violate normative expectations (great beauty, excessive wealth, etc.) A theory of inherent interest therefore enters into the system at this point.

6. Expectational Traces and Memory Retention

Predictive understanding systems naturally generate a number of expectations that are not explicitly substantiated by subsequent text. Some of these expectations form the basis for inferences, while others are discarded as "dead ends." Expectations which fail to materialize or which are contradicted by subsequent input become expectational traces in dynamic memory. Expectational traces therefore encode information about what could have happened if the story had taken a

slightly different turn. These trace structures are necessary for both answering questions and producing paraphrases.

The class of "expectational" questions query the understander about events that did not occur. [12,13] "Why didn't John stay at the party?" asks about an event which is in opposition to what happened: John left the party early. In order to answer this, an understander must have had an expectation about John staying at the party. This expectation was converted into a trace as soon as we heard that John left early. We can then search our expectational traces for the event in question, and follow the matching trace back to the event or events in the story representation that were responsible for the failed expectation.

One version of this retrieval technique was implemented in the SAM system [9,12,13] where an expectational trace was called a "ghost path." Ghost paths in SAM were entirely script-based, but in principle, traces can be derived from any predictive structure. In addition to their utility for expectational questions, expectational traces will be important in free recall as well.

Sometimes an expectational trace is "strong" enough or disturbing enough to warrant special status. Leaving a party early is not terribly disturbing, but other expectational traces should be generated with a certain reluctance. For example, in "The War of the Ghosts" we are told that the main character gets hit by an arrow, yet he feels no pain. When we are told that he was hit, we generate an expectation that the victim will undergo a negative physical state change, possibly resulting in death. But subsequent text explicitly

contradicts this expectation, giving us an expectational trace. The necessary result-causality within this particular prediction makes the trace extremely salient. It should be impossible for someone to be hit and not feel anything. This salience should result in a strong recall probability. In fact, of the ten paraphrases for this story that Bartlett published [1], eight subjects mentioned the fact that no pain was felt.

Expectational traces can be derived from a variety of knowledge structures. The ghost paths in SAM were all expectational traces derived from scriptal expectations. We can also expect to find traces resulting from goals, plans, and role themes. For example, consider the following story adapted from a newspaper account:

A small twin engine airplane carrying federal marshals and a convicted murderer who was being transported to Leavenworth crashed during an emergency landing at O'Hare Airport yesterday. During rescue operations, the murderer was able to grab a gun from a wounded marshal and effect an escape in a fire department vehicle.

In this story there is a script-based trace concerning the flight of the airplane: our airplane script had predicted a safe and uneventful journey. But consider a slight variation on this story:

A small twin engine airplane carrying federal marshals and a convicted murderer who was being transported to Leavenworth crashed during an emergency landing at O'Hare Airport yesterday. During rescue operations, the murderer was able

to grab a gun from a wounded marshal and effect an escape in a fire department vehicle.

In this story we have a new expectational trace which derives from goal-oriented predictions. The marshal should have never given the murderer his gun. On the contrary, his role theme dictates that he will try to maintain control over the murderer's movements. The marshal's goals should be in conflict with the murderer's. Given a story such as this where the marshal acts "out of character" we are forced to generate an expectational trace concerning the marshal's goals.

A taxonomy of traces is needed to determine when a trace is strong enough to warrant retrieval in free recall. Necessary causality is probably the strongest condition for recall. But it is not clear whether useful priorities can be defined in terms of purely structural factors. Will certain script-based traces tend to be stronger or weaker than plan-based traces? What are the conditions for relative strength evaluations in expectational traces? Are the conditions purely causal, or inherently determined by specific content alone? These are some of the more compelling issues surrounding expectational traces.

7. Memory Expansion Markers and Recall Errors

Errors in recall are often classified into omissions, additions, and transpositions. If we limit ourselves to those errors that omit, add, or transfigure script-based information, we can formulate these errors in terms of memory expansion markers. Memory expansion markers

point to instantiation kernels that encode all essential information necessary for a complete reconstruction of a script instantiation. For example, suppose we hear that John went to a restaurant, ordered a steak, was served by a pretty waitress, and he left a large tip. This episode will be represented by an instantiation kernel containing three types of information: (1) a script name, (2) a route through the script, and (3) specific role bindings. In this case we would have:

```
script = $restaurant
route = default path
bindings = (patron . G1) (meal . G2) (waiter/waitress . G3)
```

where G1, G2, and G3 point to memory tokens for John, a hamburger, and a pretty nameless female, respectively. If we were asked whether or not John paid a check, we would have to expand this kernel into a complete causal chain representation of the sort that the SAM system utilized [9] in order to determine that the default path of the restaurant script does, in fact, include the patron paying his check.

Many recall errors involve scriptal information and can be investigated in terms of instantiation kernels and their expansion. This particular form of dynamic memory appears to suffer greatly as the interval of retention grows. For example, in "The War of the Ghosts" there is an initial description of two young men setting out to hunt seals. In two of ten paraphrases the seals are transformed into fish, and in three of the ten there is no reference to a hunting expedition in any form. These accounts suggest that there are some standard kernel expansion problems that frequently arise and result in

memory omissions or transpositions of information.

Initially we expect that all subjects encode a memory expansion marker with the proper script name and role bindings. In this case we have a hunting script with two men bound to the role of the hunter, and seals bound to the role of prey. Early memory degradation occurs when "weak" role bindings disintegrate. The transposition of fish for seals can be accounted for in this manner. Seals have a weaker role binding than the two young men because they do not appear again in the story; the two young men are main characters. The binding for prey in the hunting script is therefore very weak and may be lost. When this occurs, and the understander attempts to reconstruct the scriptal situation, fish must be supplied by an override from associative memory. Intuitively, we expect seal hunting to be strongly linked to Eskimos, but we also expect Eskimos to fish frequently. Therefore fish can be reconstructed for the hunting episode if we know the hunters are Eskimos. To determine the precise mechanism for this is a challenging problem since it requires a strong theory of associative memory structures.

A later stage of memory degradation involves losing an entire script that appeared in the story. Possible script omission is determined by a script's causal connectivity within the larger story context: scriptal episodes that are not causally related to the rest of the story by means of enabling conditions or other causal connections are highly expendable. In "The War of the Ghosts" the initial hunting script functions only to set the scene and introduce the main characters. No subsequent causal dependencies arise from

these script activities. In the three Bartlett paraphrases which omitted the hunting script altogether, the scene was set with a simple locational description.

Free recall errors are often surprising in the ways that they defy the bounds of possibility. For example, one of Bartlett's subjects produced a paraphrase for "The War of the Ghosts" in which M1 goes home after the conversation with the warriors, and then reappears in time to accompany M2 home after the battle, as if they had never parted company. This major lapse in consistency can be explained in terms of faulty updating to a memory expansion marker. In the script instantiation kernel for the war party, M1 and M2 were originally both candidates for the role of attacking warriors. When M1 went home, the memory expansion marker should have been updated by removing M1 from this role binding. If such an update is not made, M1 will presumably go on to participate in the war party, and then return home with M2. Exactly when such updating oversights arise is a problem of some complexity, but we might expect them to occur most frequently in situations where multiple role bindings are split up after a period of shared roles. In this case, M1 and M2 had been sharing the role of hunters, and were largely indistinguishable characters up until the point where M2 was hit by the arrow.

8. Memory Configurations

We have discussed how processing at the time of understanding can generally effect tasks like question answering and paraphrase production. Each of the four dynamic memory structures are expected to affect specific recall phenomena. In addition to their isolated

affects, it is highly likely that particular configurations of dynamic memory structures will be associated with specific recall behavior. For example, consider once again "The War of the Ghosts" and our hero who felt no pain when hit by an arrow. His fellow warriors see that he was hit, and they take him back to the village. Once home, he proceeds to build a fire, and tell the villagers about the his adventures. The next morning he dies.

As we saw earlier, an expectational trace is set up by the event of being hit when our expectations about physical harm are violated. This expectation is explicitly violated when we are told he felt no pain, and then it is implicitly violated again when we hear that he went home, built a fire, and told his story as if nothing violent had happened to him. But eventually he dies and our expectation is finally substantiated. Now we have a striking memory configuration: an expectational trace was maintained for some length of time, and then was transformed into a predictive residue. The reader in this case is liable to experience an "I thought so" reaction, and we can expect the events connected to this configuration to be well preserved in memory. In general, we could predict that any transformation from an expectational trace (ET) into a predictive residue (PR) will be granted a high priority for retention, Exactly how high depends on the strength of the trace and the amount of time elapsed between trace generation and residual transformation. In our excerpt from "The War of the Ghosts" we would have roughly the following corresponding events:

Event1: Hum1 is hit by an arrow

EventA: Hum1 suffers a negative physical state change

EventB: Hum1 experiences pain

Event2: Hum1 goes home

Event3: Hum1 builds a fire

Event4: Hum1 tells his story

Event5: Hum1 dies

In our story representation these events would be organized as follows:

DYNAMIC MEMORY
CONFIGURATION:

CHRONOLOGICAL STORY
REPRESENTATION:

DERIVED CAUSAL
CONNECTIVITY:

Event1

ET (EventA, EventB, etc.)

Event2

Event3

Event4

Event5 (= EventA)

Event1



result

Event5

ET ---> PR
transformation

As we can see from this example, dynamic memory configurations of this sort are also powerful because they subsume causal relationships. We suspect that our main character died because he was hit by an arrow. This causal connection is established whenever two events are linked by a transformation from expectational trace to predictive residue. Other configurations will embed other causal links. Given this relationship between dynamic memory configurations and causal relationships, the causal connectivity analysis proposed by Schank

[23] should dovetail nicely with memory predictions based on dynamic memory representation.

Configurations which should be recognized as especially salient memory constructs include:

(1) ETs triggered by ATMs

"John was caught in an elephant stampede but he wasn't hurt."

(2) PRs triggered by ATMs

"John was caught in an elephant stampede. He was killed."

(3) expectations suspended a long time before becoming PRs

"John was caught in an elephant stampede." (After a ten page digression we are told that he was killed.)

(4) ETs transformed to PRs

M2 dying in "The War of the Ghosts"

(5) mutually exclusive PRs

"John and Bill both loved Mary. They formed a menage-a-trois."

By analyzing narratives in terms of dynamic memory configurations and collecting paraphrase data, we can expect to arrive at a comprehensive set of salient configurations. This is clearly an area where paraphrase data on reasonably long narratives is needed to further our investigation.

The issue of dynamic memory configurations is important not only for the task of paraphrase, but question answering as well. The hierarchical structure of predictive residues will be determined in

part by salient configurations involving those residues. The resulting priorities among residues will therefore reflect structural properties of the narrative's content. Because the identification of key memory configurations (such as 1-5 above) will subsume a sensitivity to content, it follows that dynamic memory configurations will reflect salient components of the memory representation and control access to those components accordingly.

9. Additional Problems

It should be clear that the representational strategy being proposed here raises a great many questions about processes at the time of understanding as well as questions about the nature of memory degradation over time. While many of these questions have been touched upon in the course of this paper, a few more deserve to be raised at this point.

From a broad perspective, we are interested in designing story representations and access strategies which will account for free and cued recall phenomena. More information about a story can be elicited in cued recall (question answering dialogues) than in free recall (paraphrase generation). We would somehow like to be able to account for this general discrepancy.

In discussing memory representations and processes of memory access, it is not always clear whether or not a specific task is being assumed. It is usually clear when someone is talking about question answering behavior; but it is far less clear when a discussion implicitly assumes free recall tasks. For example, Schank's article

"The Structure of Episodes in Memory" [23] addresses free recall access without explicitly acknowledging the fact that this is only one aspect of memory access. In this paper we are given rules like "dead-end causal chains are liable to be forgotten," and we understand from context that information in such a chain is liable to be omitted in paraphrase production. No claims are made about whether or not this information is available for processes engaged in question answering.

We must assume that for every story (of sufficient complexity) there will be some class of information in the story representation which is not accessed for paraphrase production but which can be accessed for answering questions. Access in free recall involves moving from one concept to another in memory. If a concept is not causally connected within a chain of events, it will not be attached to that chain, and may therefore be "forgotten." At the same time, a question that "cues" such a concept might lead us right to it, regardless of whether or not it was accessible from other concepts.

This cued access facility is exactly what predictive residues are designed to provide. We can think of a story representation as something that resides in memory with internal links (causal chain connectives) for free recall and external links (predictive residues) for cued recall. The internal links across concepts allow us to traverse causal chains and scenarios within the representation, while the external links allow us to jump into specific pieces of the representation on the basis of information contained in a question.

While the internal links preserve properties of causal connectivity and chronological orderings, the external links are organized in a discriminating structure based solely on the internal representation for specific events. So a story with 12 PTRANS events will have these events clustered together in the hierarchical arrangement of predictive residues regardless of their dispersal in the internally linked structure. The precise implementation of competing residues remains to be formulated, but there are two obvious strategies. (1) Predictive residues may be structurally organized in terms of a discrimination net which is searched during parsing, or (2) they may be dynamically accessed by a production-based parser. This issue of implementation will be tackled when we build a computer program that accesses dynamic memory during question answering. But beyond this question of implementation, other problems remain which are independent of the precise computational mechanisms employed.

First of all, there is the problem of accessing parts of the story representation that were not predicted at the time of understanding. Predictive residues, as we have described them, extend only to the class of information that was expected by a predictive knowledge structure. But we clearly need to be able to answer questions about events that were totally unexpected. If we hear about a singing snake in the context of a hunting trip, we should clearly be able to remember that snake, in spite of the fact that there were no predictions anticipating the snake's appearance. This situation will arise whenever an ATM is encountered which is an ATM by virtue of its unanticipated nature.

This suggests that some residues must be constructed "on the fly" in response to unanticipated information. These "spontaneous residues" will then be integrated with the predictive residues as if they were derived no differently. But we might expect to find processing limitations on the number of spontaneous residues a reader can construct. Such a limitation will manifest itself in the recall difficulties we previously attributed to "ATM overload." When too much unanticipated information occurs over an interval of processing time, the construction of spontaneous residues might be sacrificed so that the system can "keep up" with some minimal rate of input information flow. If no residues are produced for a set of ATM's, it will not be possible to access these ATM's for either question answering or paraphrase production. Alternatively, partial residues might be constructed which will allow limited access during question answering, only if the right "cue" is present in the question.

The problem of partial residues (either predictive or spontaneous) raises another interesting set of problems. Precisely what goes into a residue in the first place? Subjects frequently fail to answer questions which are underspecified when a more specific question will suffice to "jostle" their memory. For example, a subject who fails to summon information in response to "What did John tell Karen?" might then proceed to remember the interaction if next asked "Why did John lie to Karen?" Such question answering behavior can be analyzed to tell us precisely what informational components are present in a residue. At the same time, evidence for particular representational strategies can be derived from such data. For example, if a representation for "making an excuse" is structurally

similar to the representation for "telling a lie," then a residue formulated in terms of the former might be triggered by a question stated in terms of the latter. Questions about necessary and sufficient similarity conditions across representations could therefore be investigated by varying question cues and examining the resulting answers provided by a number of subjects.

The process of integrating new residues at the time of understanding also raises significant issues. We do not expect to accumulate residues into an unstructured list which will then be reorganized into hierarchical priorities after the text is completed. The process of hierarchical ordering must be completed as the narrative is processed, with old information being reordered as new information comes in. In this way, information which seems important at the time it is read may be ultimately assigned a low priority as subsequent information supercedes it. Within this framework, everything in the narrative can be present in memory, while only a limited portion of this information may be accessible. A hierarchical cut-off can occur at the N th level which renders information with priority lower than N inaccessible. This sort of mechanism is necessary when long narratives (like War and Peace) are processed. The access cutoff need not be an abrupt mandate: information at the $N+1$ st level might be accessible only if the residue is activated "very precisely" where precision is measured in terms of the number of necessary tests completed. Information at the $N+2$ nd level could then be accessed only by "totally precise" activation cues. In this way problems of variable access can be investigated in terms of hierarchical priorities on the residues and the relative strength of

activation cues in a question.

Within the framework of variable access outlined above, certain phenomena in memory degradation might be easily modelled. For example, loss of cued access over time might be merely of function of threshold raising. After one week access thresholds might move up one level; so a cutoff that was initially at the Nth level is now at the N-1st level. After two weeks access thresholds may be altered again, and so on. At some point in time we would expect the thresholds to stabilize, leaving a relatively permanent access potential.

All of the above problems require further investigation in terms of both the computational model and recall data on narratives. We have merely touched on the range of recall phenomena that can be investigated in terms of dynamic memory representations. The strategy outlined is primarily a strategy for memory access, which is itself independent of any particular scheme for conceptual representation. Within this strategy, the access of information is sensitive to the retrieval task at hand (question answering versus paraphrase), the relative importance of the information sought (its priority ranking within the hierarchy of residues), the strength of a recall cue in the case of a question asked (necessary activation conditions within a residue), and the salience of information for the purposes of paraphrase (dynamic memory configurations). Within this computational framework based on models of predictive understanding, it is possible to formulate questions about variable memory access that can be investigated in terms of both computational models and empirical data derived from human subjects.

4. Bibliography

- [1] Bartlett, F. C. (1932). Remembering: A Study in Experimental and Social Psychology. Cambridge University Press. New York.
- [2] Bobrow, D. G., Kaplan, R. M., Kay, M., Norman, D. A., Thompson, H., and Winograd, T. (1977). GUS, a frame-driven dialog system. Artificial Intelligence, Vol. 8, No. 1.
- [3] Bobrow D. G., and Norman, D. (1975). On data-limited and resource-limited processes. Cognitive Psychology. No. 7: 44-64.
- [4] Bower G. H., Black J. B., and Turner T. J. (1979). Scripts in text comprehension and memory. Cognitive Psychology. (in press).
- [5] Black, J. B., and Bower, G. H. (1979). Story understanding as problem-solving. Poetics. (in press).
- [6] Carbonell, J. G. (1979). Subjective understanding: computer models of belief systems. Research Report #150. Department of Computer Science. Yale University. New Haven, Conn.
- [7] Charniak, E. (1978). On the use of framed knowledge in language comprehension. Artificial Intelligence. Vol. 11, No. 3.
- [8] Crowder, R. G. (1976). Human Learning and Memory. Lawrence Erlbaum Assoc. Hillsdale, New Jersey.
- [9] Cullingford, R. (1978). Script application: computer understanding of newspaper stories. Research Report #116. Department of Computer Science. Yale University. New Haven, Conn.
- [10] Gershman, A. V. (1979). Knowledge-based parsing. PhD thesis. Department of Computer Science. Yale University. New Haven, Conn.
- [11] Kaplan, R. M. (1975). On process models for sentence analysis. in Explorations in Cognition. eds: Norman and Rumelhart. W. H. Freeman and Co. San Francisco.
- [12] Lehnert, W. G. (1977). Human and computational question answering. Cognitive Science. Vol. 1, No. 1.
- [13] Lehnert, W. G. (1978). The Process of Question Answering. Halsted Press. New York.
- [14] Lehnert, W. G. (1978). Representing physical objects in memory. Technical Report #131. Department of Computer Science. Yale University. New Haven, Conn.
- [15] Marcus, M. (1975). Diagnosis as a notion of grammar. Proceedings from Theoretical Issues in Natural Language Processing. Cambridge, Mass.

- [16] Minsky, M. (1975). A framework for representing knowledge. In The Psychology of Computer Vision. Ed: Winston. McGraw-Hill Inc. New York.
- [17] Norman, D. (1972). Memory, knowledge, and the answering of questions. San Diego: Center for Human Information Processing. Memo CHIP-25. University of California. San Diego, California.
- [18] Rieger, C. (1975). Conceptual memory. in Conceptual Information Processing. ed: Schank. North-Holland Pub. Co. Amsterdam.
- [19] Riesbeck, C. K. (1975). Conceptual analysis. in Conceptual Information Processing. ed: Schank. North-Holland Pub. Co. Amsterdam.
- [20] Riesbeck, C. K. and Schank, R. C. (1975). Comprehension by computer: expectation-based analysis of sentences in context. Research Report #78. Department of Computer Science. Yale University. New Haven, Conn.
- [21] Rumelhart, D. E. (1976). Notes on a schema for stories. in Representation and Understanding. (eds: Bobrow, D. G., and Collins, A.) Academic Press. New York.
- [22] Schank, R. C. (1975). Conceptual dependency. in Conceptual Information Processing. North-Holland Pub. Co. Amsterdam.
- [23] Schank, R. C. (1976). The structure of episodes in memory. in Representation and Understanding. eds: Bobrow and Collins. Academic Press. New York.
- [24] Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals and Understanding. Lawrence Erlbaum Assoc. Hillsdale, New Jersey.
- [25] Schank, R. C. and Carbonell, J. C. (1978). Re: The Gettysburg Address: representing social and political acts. Research Report #127. Department of Computer Science. Yale University. New Haven, Conn.
- [26] Simmons, R. F. (1979). Towards a computational theory of textual discourse. (to appear in) Journal of Computational Linguistics.
- [27] Wilensky, R. (1978). Understanding Goal-Based Stories. Technical Report #140. Department of Computer Science. Yale University. New Haven, Conn.
- [28] Wilks, Y. (1973). An artificial intelligence approach to machine translation. In Computer Models of Thought and Language. Eds: Schank and Colby. W. H. Freeman and Co. San Francisco.
- [29] Winograd, T. (1972). Understanding Natural Language. Academic Press. New York.

[30] woods, W. A. (1969) Augmented transition networks for natural language analysis. Rep. CS-1, Computer Laboratory, Harvard University. Cambridge, Mass.

[31] Woods, W. A., Kaplan, R. M., and Nash-Webber, B. (1972). The lunar sciences natural language information system: final report. BBN Report No. 2378. Bolt, Beranek and Newman, Inc. Cambridge, Mass.